

How do the dynamics of battery discharge affect sensor lifetime?

Laura Marie Feeney
Swedish Institute of Computer Science
Kista, Sweden
Contact: lmfeeney@sics.se

Christian Rohner
Per Gunningberg
Uppsala University
Uppsala, Sweden

Anders Lindgren
Lars Andersson
Pricer AB
Stockholm, Sweden

Abstract—Evaluation of energy consumption and device lifetime in battery-powered wireless sensor networks (WSN) is almost exclusively based on estimates of the total charge (i.e. mA-h) consumed by the device. In reality, batteries are complex electro-chemical systems and their discharge behavior depends heavily on the timing and intensity of the applied load. However, there is very little empirical data or reliable models available for the kinds of batteries and loads that are typically used in WSN. The effect of battery dynamics on sensor lifetime is therefore not well understood.

We characterize CR2032 Li coin cells using carefully controlled synthetic loads and a wide range of WSN-typical load parameters. Our results are the first to quantify in-depth the discharge behavior of primary batteries in the WSN context. We report that in some common cases, observed lifetimes can differ from predicted ones by almost a factor of three. Furthermore, loads with similar average currents – which would be expected to have similar lifetimes – can vary significantly in the amount of capacity they can utilize, with short duration loads generally faring better.

The results show that energy evaluation based on a mA-h consumed model has significant limitations. This has important implications for the design and evaluation of WSN applications, as well as for practical problems in network dimensioning and lifetime prediction.

I. INTRODUCTION

Battery-constrained wireless sensor networks (WSN) are designed to maximize their useful lifetime by minimizing the battery capacity (mA-h) consumed by devices over time. Estimates of this value are the basis of almost all WSN energy performance evaluation and lifetime prediction.

Under this energy consumption model, the battery itself is treated as a simple store of charge. In reality, batteries are complex electro-chemical systems: The timing and intensity of the applied load determines how much of the battery's nominal capacity can be utilized before the output voltage drops below the level needed to operate the device.

The macroscopic properties of battery discharge – rate dependent capacity, charge recovery, higher sensitivity to load at low SoC – are well known. However, quantitative results are highly specific to each battery chemistry and structure, as well as to the operating regime. Very little data is available for the kinds of inexpensive, primary (non-rechargeable) batteries and high current/short duration loads that are found in many WSN applications.

The impact of battery dynamics on sensor lifetime is therefore a poorly understood aspect of WSN performance. Without data about battery output voltage in response to load, it is not clear whether the “mA-h consumed” model provides a sufficiently accurate view of device lifetime. Clarifying this has important implications not only for design and evaluation of WSN hardware and software, but also for dimensioning, lifetime prediction, and load balancing in deployed networks.

The contributions of this work are as follows:

- We present a large scale characterization of the Panasonic CR2032 battery, an inexpensive, primary lithium coin cell that is often used in body-area WSN. Our results are the first to thoroughly quantify battery discharge behavior in the WSN context. The measurements were obtained using a custom testbed that is capable of generating carefully controlled synthetic loads using a wide range of WSN-typical parameters. Over 50 systematically defined combinations of load parameters were examined.
- Our data show that at low duty cycles, the observed lifetime differs from the lifetime predicted by “mA-h consumed” models by *as much as 260%*, even when the effect of high loads on capacity is taken into account.
- We examine sets of periodic loads with the same time-average current, but different load values, duty cycles, and active periods to explore the impact of various load parameters on sensor lifetime. Such loads are expected to have the same lifetime, since charge is being consumed at the same rate. Instead, we observe differences in lifetime of *as much as 15-20%*. The data suggest that shorter load durations are associated with longer observed lifetimes, especially high current loads at low duty cycles.

These results quantify the approximation associated with the ubiquitous “mA-h consumed” model of battery consumption, with implications for evaluating WSN energy performance and estimating device lifetime. The data also provide a solid empirical basis for future work in truly “battery-aware” methods in WSN applications.

II. BACKGROUND

Our test batteries were CR2032 lithium coin cells, with nominal 3.0V output voltage and 225 mA-h capacity (at rated capacity of 200 μ A), manufactured by Panasonic. Because of

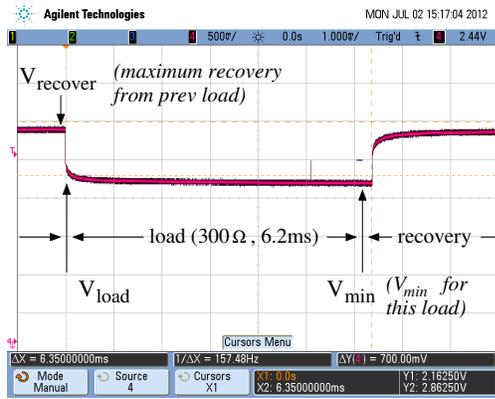


Fig. 1. CR2032 Li-coin cell: output voltage in response to a 300Ω (10mA) load with 6.2 ms duration.

its low cost and small size, the CR2032 is especially popular for sport and body-area WSN applications, as well as for computational RFID and logistics applications. The slightly larger CR2354 battery (560 mA-h), which is also used in animal monitoring networks and smaller general purpose WSN platforms, has the same chemistry and similar structure.

A battery consists of an anode and cathode, separated by electrolyte and a permeable membrane (to prevent an internal short). In Li-MnO₂-based lithium batteries (CRxx), the Li anode is oxidized ($\text{Li} \rightarrow \text{Li}^+ + \text{e}^-$) and the Li⁺ ions diffuse through the electrolyte (an organic solvent) – the corresponding electrons flow through the applied load – to the MnO₂ cathode, which is reduced. This reaction is not reversible, so this is a non-rechargeable battery¹.

As the battery is discharged and the active species at the anode and cathode are consumed, the battery’s state of charge (SoC) decreases. This means not only that the residual charge capacity has decreased, but also that there have been changes in the chemical and physical properties of the electrodes and the composition of the electrolyte. These changes affect the battery’s electrical properties, especially the internal resistance and the efficiency of the electro-chemical reactions. Moreover, these effects are highly specific to each battery chemistry and structure, and (to a lesser extent) even to a particular manufacturer. It is therefore impossible to generalize quantitative results across different types of batteries. The details of these processes are the domain of material chemistry and far beyond the scope of this work. For our purposes, we need only recognize that battery voltage has complex dependencies on the timing and intensity of the applied load and focus our attention on macroscopic behaviors.

Figure 1 shows an oscilloscope trace of the output voltage of a CR2032 Li-coin battery in response to load. When the load is applied, there is an immediate drop in output voltage (V_{load}), caused by the battery’s internal resistance. The output voltage then continues to decrease slightly over the duration of the load (V_{min}). This voltage drop reflects how well the electro-

¹Note that the chemistry and structure of secondary (rechargeable) Li-ion batteries is quite different.

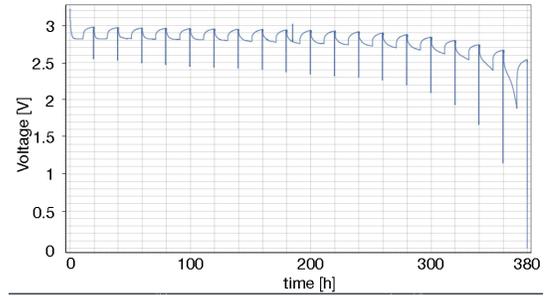


Fig. 2. CR2032 Li-coin cell: voltage response over a sequence of periodic loads (ETC Battery and FuelCells Sweden AB).

chemical reactions in the battery “keep up” with the demand from the applied load, reflecting the accessibility of active species at the electrodes and efficiency of transport through the electrolyte. The voltage partially recovers after the load is removed, with the immediate recovery from the battery’s internal resistance and the slower one reflecting processes that relax the changes in the electro-chemical state that have accumulated during discharge (V_{recover}).

Figure 2 shows the output voltage in response to a periodic load over a longer time scale. The load is composed of two alternating elements: A 1mA load is applied for 11h, followed by a rest time of 8.5h – each active period reduces the state of charge (SoC) by ~5%. A 22mA load is applied for 10s, followed by a 30min rest, demonstrating how the load response varies with the SoC. (This load is based on a standard battery test sequence.) The load and recovery effects are clearly visible: The high current load causes a much larger voltage drop than the low current load. The internal resistance increases as SoC decreases, due to electro-chemical changes that impede the flow of ions. As the battery SoC decreases, the recovery effect also becomes smaller and the voltage drop over the duration of the load also becomes much steeper, due to depletion of active species and other changes in the battery chemistry and structure.

Eventually the output voltage falls below some cut-off voltage, which is determined by the requirements of the device using the battery (usually between 1.8V and 2.2V). Even though there may be considerable charge remaining in the battery, it cannot be provided at sufficiently high voltage to operate the device correctly. It is this cut-off voltage that actually determines the device lifetime.

At the macroscopic level, we abstract these complex phenomena as rate dependent capacity and charge recovery. Rate dependent capacity refers to the fact that discharging the battery at high current results in large voltage drops that reduce the amount of capacity that can be used before reaching the cut-off voltage. Charge recovery refers to the fact that intermittent discharge, with intervals of low or zero load that allow the output voltage to recover, increases the amount of capacity that can be used.

Temperature is another key factor in battery performance, due to its impact on the electro-chemical processes described



Fig. 3. Test card (Rev 2): The processor (bottom right) connects load resistors to each battery and records output voltage. Cost is <math><150\text{€}</math>.

above. Our experiments were all performed at room temperature and we defer this very important topic to future work.

III. EXPERIMENT METHODOLOGY

Our goal is to quantify key battery discharge behaviors and understand what effects are most significant for WSN development. The experiment methodology was therefore driven by the need to run large numbers of experiments for systematic exploration of load parameters. This also meant it was more important to be able to vary the load current, duration, and duty cycle over a range of realistic values than to reproduce the load generated by specific sensor hardware or protocols.

In previous work, batteries were discharged by connecting them to operating sensor hardware, which leads to poorly controlled loads. An operation may involve multiple device components or state transitions and the cost of operations such as frame transmission can be variable. Furthermore, the load can only be controlled indirectly, via system parameters like transmit power or packet length. To avoid these issues, we designed the testbed to generate controlled synthetic loads.

In addition, because WSN applications are intended to have low loads, realistic discharge times are long (weeks). For practical reasons, it was therefore important to be able to run many experiments in parallel, which also placed some constraints on cost per experiment.

A. Testbed

We developed custom hardware (Figure 3) that allows us to apply a resistive load to sets of batteries according to a user-defined schedule and measure the battery output voltage. The boards are connected via USB to a user PC, which sends configuration commands to the boards and collects data from them. The testbed is described in more detail in [1].

Each board has an ATmega16U4-AU controller and eight battery holders, each connected to a set of four swappable resistors. The controller can be programmed to connect and disconnect any combination of resistors to the batteries and to record voltage measurements via its ADC interface.

The control firmware supports user-defined continuous and periodic loads. The loads used in this work are essentially simple square waves. However, because the loads are resistors rather than current-controlled elements, the load current decreases over time, along with the battery output voltage.

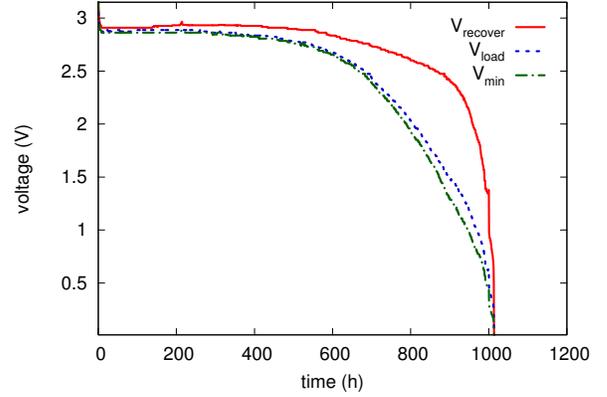


Fig. 4. Raw data for a single battery: $750\ \Omega$ (4mA) load with duration $75\text{ms}/1.0\text{s}$ (7.5% duty cycle).

Loads can be scheduled with sub-millisecond granularity and validated using an oscilloscope.

Commercial battery test equipment provides more fine grain control, including current and voltage controlled loads and more precise measurements. Such equipment is expensive: A commercial system costs over $\text{€}35000$ and has less than one-third of the capacity of our 20 board (160 battery) testbed, which cost less than $\text{€}3000$. There are also companies that do battery testing for safety certification (the data in Figure 2 is from such a service). For more than a few batteries, this approach is also prohibitively expensive. With identification of macroscopic effects as the goal, the more scalable approach was preferred.

B. An experiment in detail

To illustrate how the measurement data is processed and interpreted, we examine a representative experiment in detail. The experiment used a $750\ \Omega$ (4mA) load with a duration of 75ms and period of 1 second. In practical terms, these values suggest an RFM TR1001 radio receiving one ~ 150 -byte frame per second. The load has a duty cycle of 7.5% and a nominal average current of $300\ \mu\text{A}$ ($4\text{mA} \times 7.5\%$). Because the battery voltage decreases over time, the load (V/R) also decreases over time. This means that the actual average current is less than the nominal value: $\sim 276\ \mu\text{A}$. (In the text, currents are referred to using their nominal value.)

The testbed hardware measures the battery output voltage at user defined intervals. For periodic loads, the voltage is measured at each of the three key points shown in Figure 1. These are V_{recover} , measured just before the load is applied; V_{load} , measured just after the load is applied; and V_{min} , measured just before the load is removed and assumed to be the minimum voltage over the period. Since a device fails when the battery cannot maintain the required output voltage under load, we take V_{min} as the determinant of device lifetime. (In the following text, voltage refers to V_{min} .)

Over the 1000 hour duration of the experiment, nearly 4 million loads were measured for each battery. We first examine the raw data obtained from one battery. Figure 4 shows the

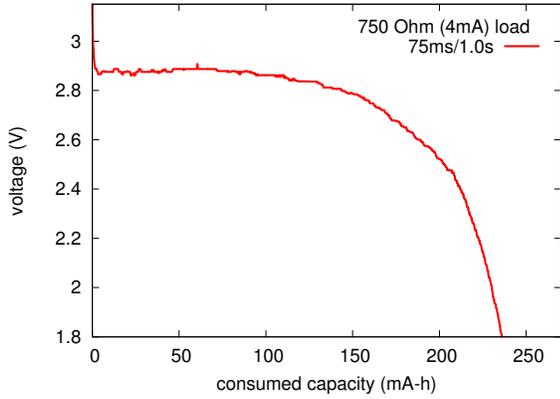


Fig. 5. Output voltage vs the consumed capacity. When the output voltage reaches 2.0V, about 229 mA-h of charge has been extracted from the battery.

values of V_{recover} , V_{load} , and V_{min} (y-axis) associated with a load applied at time t (x-axis) and how they evolve over time. V_{min} first drops below 2.0V at about $t=800\text{h}$, which would be the lifetime for a device with that cut-off voltage.

When comparing experiments that have different average currents, it is often useful to consider the capacity that has been consumed when the cut-off voltage is reached, rather than time that it takes (i.e. normalizing with respect to current). This allows us to highlight differences in lifetime that reflect how load timing and intensity affect the usable capacity, rather than the obvious difference that comes from consuming capacity at different rates. Since even loads with the same nominal average current will have slightly different actual currents, most results are reported this way. Figure 5 shows the output voltage (V_{min}) from Figure 4 plotted vs consumed capacity, rather than time.

There is variation among batteries and also premature battery failure, discussed further in Section IV. When averaging over a set of batteries from a given experiment, we use the mean of the $n = 3$ batteries with the largest V_{min} values. The point is not to exaggerate the absolute capacity, but rather to focus on the behaviors that arise from the electro-chemical properties of the battery, rather than manufacturing variation.

C. Experiments

Consider the three idealized loads shown in Figure 6. They all have the same time-average current and consume battery capacity (integral of current over time) at the same rate. These loads would therefore be expected to have the same lifetime under both simple “mA-h consumed” models and models that take rate dependent capacity into account. Differences in lifetime between these loads can therefore be used to study the effects of load parameters. The load currents and duty cycles, as well as the absolute load durations and periods used in these experiments are structured around this principle (Table I).

We use loads ranging from $15\text{k}\Omega$ to 120Ω , with most periodic loads using 4-25 mA nominal currents. For reference, a CC2420 IEEE 802.15.4 transceiver[2] consumes 8.5-17.4mA for transmitting and 18.8mA for receiving. The low data

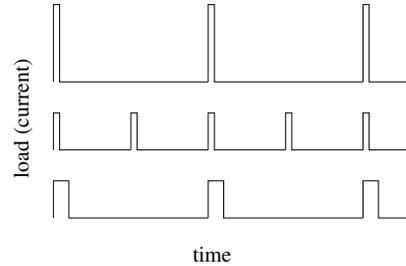


Fig. 6. Three idealized loads with the same average current. The load parameters vary systematically: e.g. cutting both the load and the period in half or cutting the load in half and doubling the duration.

rate RFM TR1001[3] transceiver consumes up to 12mA for transmitting and 1.8-3.8mA for receiving.

We concentrated on periodic loads with actual time-average currents of $250\text{-}800\mu\text{A}$ and lifetimes of 10-40 days. Duty cycles ranged from 1.2% for high currents to 30% for the lower ones. We also ran some experiments with extremely high duty cycles (50-90%) to compare with results from earlier work.

Load durations ranged from 2.4ms to 150ms. These are fairly typical times for operations such as channel sensing, wakeup preambles, and transmitting short frames using high (250kbps) and low (19.2kbps) bit rate transmissions.

Load periods ranged from 20ms to 2s. These are typical of the periods used in various WSN MAC protocols and sensing applications. These values also span the range of “optimal” wakeup schedules that were derived in [4] for several MAC protocols, subject to various delay and reliability constraints.

A 2.0V cut-off voltage was used for all of the results reported here. For reference, the CC2420 IEEE 802.15.4 transceiver has a minimum input voltage of 2.1V (without voltage regulator) and the RFM TR1001 transceiver has a minimum input voltage of 2.2V.

IV. EXPERIMENTAL RESULTS

This section presents four sets of results: We measure lifetime under various continuous loads to determine the rate dependent capacities. Then we use these capacities to estimate lifetime for periodic loads at various duty cycles. Then we measure the capacity consumed by a variety of periodic loads, focusing on differences between loads with similar time-average currents. Finally, we present statistics on battery variation within experiments.

A. Continuous discharge and rate dependent capacity

Measurements of the battery output voltage under continuous load are used to evaluate the rate dependent capacity. Figure 7 shows how the battery output voltage decreases as the battery capacity is consumed. Not only does the higher load current drain the available capacity more quickly, it is also only able to use a smaller portion of the capacity before reaching a cut-off voltage. The figure demonstrates the latter effect, rather than the factor of n lifetime difference between currents. (This is the normalization mentioned in the previous section.)

		average current (nominal) and duty cycle						
load ↓	duration ↓	0.3mA	0.6mA	0.9mA	1.2mA	3.0mA	6.0mA	7.5mA
750 Ohm (4mA) duration and → periods	15	7.5%	15%	22.5%	30%	75%	–	–
	30	200	100	66.7	50	20		
	45		200					
	60			200				
	150	2s	1s	666.7	200	500	200	
300 Ohm (10mA)		3.0%	6.0%	9.0%	12%	30%	60%	75%
duration and → periods	3 ms	100 ms						
	6	200						
	15			166.7		50	25	20
	18			200				
	30	1s				500	250	200
150								
120 Ohm (25mA)		1.2%	2.4%	3.6%	4.8%	12%	24%	30%
duration and → periods	2.4 ms	200 ms						
	7.2	600		200				
	12	1s						
	15			416.7				50
	150							500

TABLE I

EXPERIMENTS (SUBSET): THE LEFTMOST COLUMN IS THE LOAD VALUES. THE NEXT COLUMN IS A SET OF LOAD DURATIONS (MS). FOR EACH LOAD DURATION, THE CORRESPONDING ROW LISTS THE PERIOD (MS EXCEPT WHERE NOTED) AT WHICH THE LOAD IS REPEATED. EACH COLUMN IS A SET OF EXPERIMENTS WITH THE SAME NOMINAL AVERAGE CURRENT. OVER 50 EXPERIMENTS WERE RUN, WITH 8-10 BATTERIES EACH.

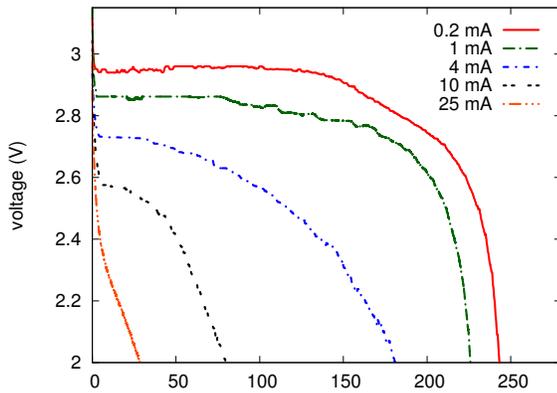


Fig. 7. Rate dependent capacity: A higher load extracts less capacity for all cut-off voltages. Not only does a 4mA load drain the battery at a rate 4x times faster than the 1mA load, but the available capacity is also ~20% less.

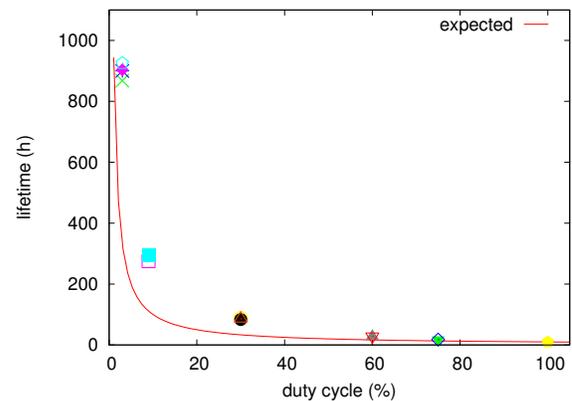


Fig. 8. Device lifetime (for 2.0V cut-off) for a 10mA load at different duty cycles. The solid line shows the expected lifetime, based on the lifetime at 100% duty cycle. The data points are average observed lifetimes for 10mA loads with the given duty cycle and different absolute periods.

The measured capacity for the 15kΩ (200 μA) reference load (243 mA-h) is slightly higher than the nominal capacity of 225 mA-h given by Panasonic[5]. Similarly at 1mA, the measured capacity was 220mA (vs 200mA). (No data is given for higher loads.) Presumably the specification is conservative, to compensate for variation among cheaply manufactured batteries.

B. Lifetime estimation at various duty cycles

The solid line in Figure 8 shows the estimated lifetime (to 2.0V cut-off) for a variety of 10mA loads with different duty cycles, using the rate dependent capacity under a continuous 10mA load (100% duty cycle) from the experiments of Figure 7. Each data point represents the observed lifetime (averaged

over max-n batteries, as described in Section III) for a different combination of load parameters at the given duty cycle (e.g. 15ms/200ms vs 150ms/2s),

At high duty cycles, the observed and estimated lifetimes are similar. For lower, more realistic duty cycles, the observed lifetimes are significantly larger. The difference between high and low duty cycles highlights the risk of extrapolating results from short experiments run at unrealistically high loads.

Figure 9 presents this data as the relative difference between the observed and expected lifetime, for a wider range of loads. Once the duty cycle is below ~60% (4mA loads) or ~30% (10mA loads), the observed lifetime seems to stabilize at 1.2x and 2.6x times larger than would be expected based on the

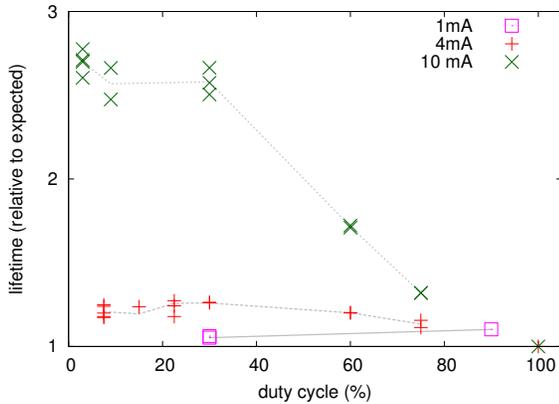


Fig. 9. Observed lifetime relative to expected lifetime based on rate dependent capacity for the load current. At a given duty cycle, the observed lifetimes vary somewhat depending on absolute load duration and period (vertical clusters of data points).

smaller duty cycle alone.

This suggests that there is a point at which further reducing the duty cycle does not result in additional benefits from charge recovery. Since higher loads put more stress on the battery, the gain is larger and continues to accrue over a larger range of duty cycles than it does for lower loads.

We also note some variation among observed lifetimes at the *same* load and duty cycle. These experiments differ only in the absolute load duration and period, so they would be expected to have the same lifetime. The differences between them, seen in the vertical spread of data points at each duty cycle, suggests that load duration also plays a role in determining capacity and lifetime.

C. Parameter exploration

The next sequence of results explores the relationship between a load's time-average current and the battery capacity it utilizes before reaching a cut-off voltage (2.0V). Systematic variations in load parameters were used to define sets of experiments with different load durations, periods, and currents, but the same time-average current (columns in Table I). These experiments allow us to investigate how capacity depends on average current (i.e. another view of rate dependent capacity) and also the extent to which capacity depends on the timing and intensity of the load.

The data are shown in Figures 10 and 11: The specified capacity (horizontal line) is 225 mA-h [5]. The rate dependent capacity (slanted line) is based on capacities measured using a continuous current. Compared to these two references, each data point represents the results of an experiment using a different combination of periodic load parameters from Table I and shows the mean capacity extracted from the battery vs the time average current for the load.

The rate dependent capacity effect is clearly visible: As the average current increases, the amount of charge that can be extracted from the battery decreases. The relationship observed in the data points from measurements of periodic

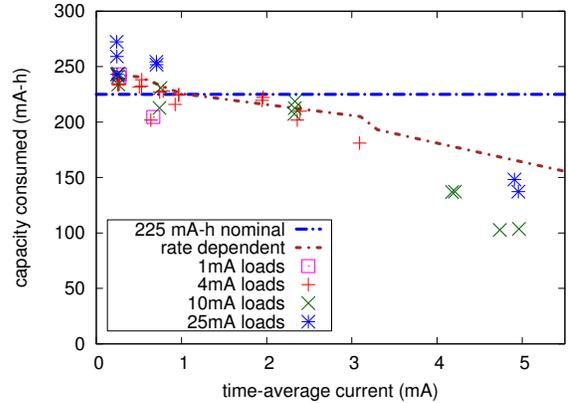


Fig. 10. Capacity vs average current: Each data point represents a different load configuration specified in Table I. Loads with the same peak current

loads is roughly similar to the data based on measurements of continuous loads with the same average current. This method provides better lifetime estimates than those based on the rate dependent capacity of the peak load, used in the previous section (Figures 8 and 9).

However, the difference in capacity consumed by loads with very similar average current cannot be explained by differences in rate dependent capacity. Some other aspect of load timing and intensity must be at work. For example, we see that a load with average current of $239\mu\text{A}$ consumes 271 mA-h capacity, while one with an average current of $251\mu\text{A}$ current consumes 234mA-h. This $12\mu\text{A}$ (4.8%) difference in average current results in a $\sim 15\%$ difference in consumed capacity and $\sim 19\%$ (47 vs 39 days) difference in lifetime for the two loads.

This result quantifies the amount of approximation associated with using energy performance metrics based on total mA-h consumed. Even taking rate dependent capacity into account (which many widely used methods do not), it is not possible to predict this rather significant difference in lifetime.

Figure 11 shows detail of Figure 10, highlighting the low time-average currents that are most relevant to WSN. For a given average current, higher peak currents (i.e. lower duty cycles) have better capacity utilization. Although this seems to conflict with the general principle of rate dependent capacity, Figure 2 suggests an explanation. As the battery SoC decreases, the voltage drop that occurs over the duration of the load becomes relatively larger and steeper. A higher load has a higher initial voltage drop, but a shorter load duration means that there is less time for the secondary process to have an impact.

D. Battery variation

There is also some variation observed between individual batteries subject to the same load. Figure 12 shows output voltage vs time for all the batteries measured in the experiment of Section II. More generally, for the $300\mu\text{A}$ group of experiments, the measured standard deviation was 14-29h on lifetimes of 800-900h. We also note that $\sim 2\text{-}3\%$ of the batteries we measured (out of many hundreds) exhibited early

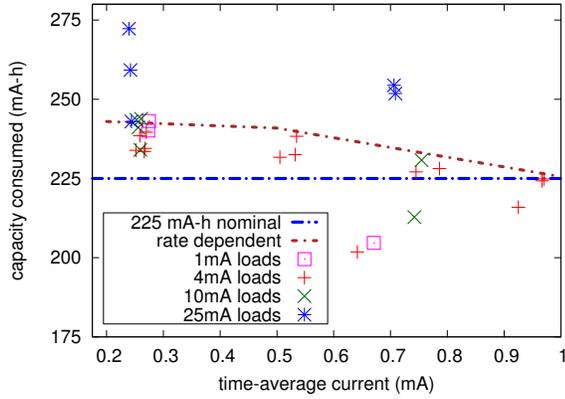


Fig. 11. Detail of Figure 10.

failure or erratic output voltage after an initial period of apparently normal operation.

This is not unexpected, given the low precision manufacturing of inexpensive batteries, but is of practical importance to WSN developers. For industrial consumers of these batteries, one of the most important factors that distinguishes among various manufactures is their consistency.

V. RELATED WORK

A. The WSN context

There are three main techniques for evaluating energy consumption in WSN. Specialized hardware can be used to measure the current drawn from the battery, e.g. [6], [7]. More commonly, system software is instrumented to record the time spent performing various operations, e.g. [8] and this trace is combined with information about the cost of each operation to determine the consumed capacity. This approach is evaluated in [9], which concludes that high accuracy can be obtained by calibrating the system with careful measurements of operation cost. WSN simulation, e.g. [10], [11], [12] can similarly generate detailed traces of the operation of the simulated device. All of these methods only consider the capacity consumed by the device and treat the battery as a simple “bucket of mA-h”.

There have been very few studies of the small, inexpensive batteries typically used in WSN. Although rate dependent capacity and charge recovery effects can be seen in earlier results, no previous work has used carefully controlled loads and a broad range of realistic parameters to systematically explore battery discharge in the WSN operating regime.

In [13], the authors measured the CR2354 coin cell (similar to the CR2032 used in our experiments). The load was generated by an AVR/RFM DR3000 transceiver, which was powered via a pulse frequency modulation DC-DC converter. The sensor’s current draw (4-12mA) is therefore seen at the battery as a high frequency load with a very high peak current (80mA). The authors highlight that DC-DC efficiency is only about 70% (such converters are not generally used in modern sensors). Although such loads are no longer representative

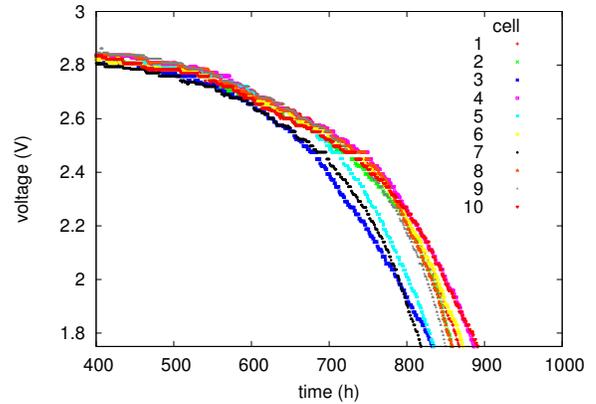


Fig. 12. Ten batteries discharged using the same load (Figure4): The key observation is the considerable variation among the discharge curves, rather than the time evolution of any specific battery.

of sensor hardware and the authors do not attempt to draw quantitative conclusions, rate dependent capacity and charge recovery effects can be seen in the data.

In [14], the authors also studied the CR2354. Measurement data were used to partially parameterize the DUALFOIL electro-chemical simulator, which is actually intended for a very different rechargeable Li-ion battery chemistry. The authors modified an (unspecified) subset of the simulator’s >50 chemical and physical parameters, until its output resembled their experimental data, though the accuracy of this approach was not reported. The authors then built simulation and emulation tools to model sensor lifetime for the Mica2DOT sensor and CC1000 transceiver. Significant rate dependent capacity and charge recovery effects can be seen in the results, although they are not systematically quantified or compared with direct measurements. The duty cycles studied in this work were also unrealistically high (25-80%).

In [15], device lifetime is determined by intentionally running batteries to depletion in a testbed. The authors measured device performance while Telos-B nodes (using AA batteries) performed a complex task sequence including duty-cycling, transmission, reception, and logging to flash memory. The authors report significant variation between battery brands, as well as differences between observed lifetime and that predicted by a simple mA-h consumed model. Unlike our results, the observed lifetimes are shorter than the predicted ones as the duty cycle drops below 75% toward 25%. However, rate dependent capacity does not seem to have been taken into account and again the duty cycle is unrealistically high. Despite interesting observations, the complex load and many interacting systems make it difficult to isolate and understand battery related effects.

Finally, Nordic Semiconductor and Energizer recently published a report [16] describing the evaluation of pulsed discharge patterns for their CR2032 batteries. This is similar to our work, but uses much higher loads (which the battery was shown to tolerate) and only a few sets of parameterizations.

B. Battery modeling

Battery modeling was first introduced to the mobile computing community in the late 1990's due to growing interest in devices like PDA's. There are many approaches to battery modeling [17], ranging from abstract analytic models to detailed electro-chemical simulations. The former must be parameterized and tuned using experimental data, while the latter are highly battery specific, requiring dozens of parameters to describe its chemical and structural properties. An older analytic model introduced in [18] for rechargeable Li-ion batteries was recently ported to ns-3 [12].

In general, there seems to have been relatively little interest in modeling the small, inexpensive primary batteries intended for WSN applications. However, an electro-chemical model [19] for Li-MnO₂ Li coin cells has recently become available [20]. An evaluation [21] of this simulator and two other abstract battery models parameterized using data from our testbed has proved somewhat equivocal.

VI. CONCLUSION AND FUTURE WORK

This paper has posed a novel question of practical importance to the WSN community: *How do the dynamics of battery discharge affect sensor lifetime?*

A battery's output voltage is determined by complex electro-chemical processes that depend on the timing and intensity of the load during discharge. The output voltage determines sensor lifetime: When the battery is no longer able to maintain a sufficiently high output voltage in response to the load presented by the device, it will fail to operate correctly.

This work is the first to provide an in-depth quantitative description of battery discharge behavior in the WSN context. We have developed a large scale testbed that is capable of generating carefully controlled synthetic loads and used it to systematically characterize the Panasonic CR2032 non-rechargeable lithium coin cell under a wide range of WSN-typical load parameters.

We have focused on macroscopic properties such as rate dependent capacity and charge recovery and their impact on sensor lifetime. Our results show that there can be errors of almost a factor of 3 in lifetime prediction, when modeling low duty cycles. Furthermore, even loads with very similar average currents can have differences of up to 15-20% in device lifetime. In general, higher loads (and hence lower duty cycles) are associated with longer observed lifetimes.

In the longer term, these results will contribute to improved methods for battery-aware design and evaluation of WSN protocols and systems. Options for future development range from using empirical models to add correction factors to "mAh consumed" lifetime estimates, to developing analytic models of the discharge process that can be parameterized using measurement data and integrated with existing simulation tools and operating counting techniques that already model the load current generated by the device. Because WSN loads are more complex than the simple square waves used in our study, an obvious next step will be to characterize the behavior of more realistic combined loads.

Another practical problem is estimating battery state-of-charge (or lifetime), especially in real time and with limited WSN resources. This is needed for network dimensioning, to provide advance warning of coverage failures in deployed networks, and for battery-aware load balancing. Battery state is affected not only by the cost of a sensor's prescribed operations, but also by external factors like interference conditions, temperature (key future work), and battery variation. A combination of improved battery models and empirical discharge data may enable new SoC estimation techniques.

REFERENCES

- [1] L. M. Feeney, L. Andersson, A. Lindgren, S. Starborg, and A. Ahlberg Tidblad, "Poster abstract: A testbed for measuring battery discharge behavior," in *7th ACM Int'l Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH)*, 2012.
- [2] "Texas Instruments," <http://www.ti.com/product/cc2420>.
- [3] "RF Monolithics, Inc." <http://www.rfm.com>.
- [4] M. Zimmerling, F. Ferrari, L. Mottola, T. Voigt, and L. Thiele, "ptunes: Runtime parameter adaptation for low-power mac protocols," in *11th Int'l Conf. on Information Processing in Sensor Networks (IPSN)*, 2012.
- [5] "Panasonic," <http://www.panasonic.com>.
- [6] P. Dutta, M. Feldmeier, J. Paradiso, and D. Culler, "Energy metering for free: Augmenting switching regulators for real-time monitoring," in *Int'l Conf. on Information Processing in Sensor Networks (IPSN)*, 2008.
- [7] A. Hergenroder, J. Wilke, and D. Meier, "Distributed Energy Measurements in WSN Testbeds with a Sensor Node Management Device (SNMD)," in *Workshop Proc of the 23d Int'l Conf on Architecture of Computing Systems*, 2010.
- [8] A. Dunkels, J. Eriksson, N. Finne, and N. Tsiftes, "Powertrace: Network-level power profiling for low-power wireless networks," Swedish Institute of Computer Science, Tech. Rep. T2011:5, 2011.
- [9] P. Hurni, B. Nyffenegger, T. Braun, and A. Hergenroeder, "On the accuracy of software-based energy estimation techniques," in *8th European Conf. on Wireless Sensor Networks (EWSN)*, 2011.
- [10] V. Shnayder, M. Hempstead, B.-R. Chen, G. W. Allen, and M. Welsh, "Simulating the power consumption of large-scale sensor network applications," in *2nd Int'l Conf. on Embedded Networked Sensor Systems (SenSys)*, 2004.
- [11] L. M. Feeney and D. Willkomm, "Energy framework: An extensible framework for simulating battery consumption in wireless networks," in *3rd Int'l Workshop on OMNeT++*, 2010.
- [12] H. Wu, S. Nabar, and R. Poovendran, "An energy framework for the network simulator 3 (ns-3)," in *4th Int'l ICST Conf. on Simulation Tools and Techniques (SimuTools)*, 2011.
- [13] S. Park, A. Savvides, and M. B. Srivastava, "Battery capacity measurement and analysis using lithium coin cell battery," in *Int'l Symp. Low Power Electronics & Design (ISLPED)*, 2001.
- [14] C. Park, K. Lahiri, and A. Raghunathan, "Battery discharge characteristics of wireless sensor nodes: An experimental analysis," in *IEEE Conf. on Sensor and Ad-hoc Communications and Networks (SECON)*, 2005.
- [15] H. A. Nguyen, A. Förster, D. Puccinelli, and S. Giordano, "Sensor node lifetime: an experimental study," in *IEEE Int'l Conf. on Pervasive Computing and Communications Workshops (PEROCM)*, 2011.
- [16] K. Furset and P. Hoffman, *High pulse drain impact on CR2032 coin cell battery capacity*, Nordic Semiconductor and Energizer, Sept 2011, Technical memo.
- [17] R. Rao, S. Vrudhula, and D. Rakhmatov, "Battery modeling for energy aware system design," *IEEE Computer*, vol. 36, no. 12, 2003.
- [18] D. Rakhmatov, S. Vrudhula, and D. Wallach, "A model for battery lifetime analysis for organizing applications on a pocket computer," *IEEE Trans. on Very Large Scale Integration (VLSI) Systems*, vol. 11, no. 6, 2003.
- [19] G. Yeduvaka, R. Spotnitz, and K. Gering, "Macro-homogenous modeling of commercial, primary Li/MnO₂ coin cells," *ECS (Electrochemical Society) Trans.*, vol. 19, no. 16, 2009.
- [20] "Battery Design Studio," 2012, <http://www.batsdesign.com>.
- [21] C. Rohner, L. M. Feeney, and P. Gunningberg, "Evaluating battery models in wireless sensor networks," in *11th Int'l Conf. on Wired/Wireless Internet Communications (WWIC 2013)*, 2013.